

## Condition Variables

- multithreaded programs need more than just locks
- locks support mutual exclusion
- *condition variables* support waiting for a condition to become true
  - e.g. for a child thread to end (join)
  - e.g. producer/consumer (bounded buffer)
  - e.g. covering conditions

## The Need For Condition Variables

- as we've seen with locks, spinning for any length of time is inefficient

```
volatile int done = 0;

void *child(void *arg) {
    printf("child\n");
    done = 1;
    return NULL;
}

int main(int argc, char *argv[]) {
    printf("parent: begin\n");
    pthread_t c;
    pthread_create(&c, NULL, child, NULL); // child
    while (done == 0)
        ; // spin
    printf("parent: end\n");
    return 0;
}
```

## Condition Variables – pthreads API `#include <pthread.h>`

- a condition variable is just a queue for waiting threads

- declaration, initialization, and cleanup

```
int pthread_cond_init(pthread_cond_t *cond,
                    const pthread_condattr_t *attr);
int pthread_cond_destroy(pthread_cond_t *cond);
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

- PTHREAD\_COND\_INITIALIZER initializes with default values

- put current thread to sleep (*wait*)

```
int pthread_cond_wait(pthread_cond_t *cond, pthread_mutex_t *mutex);
int pthread_cond_timedwait(pthread_cond_t *cond,
                          pthread_mutex_t *mutex, const struct timespec *abstime);
```

- wait requires first holding a lock (*mutex*)

- wake up one (*signal*) or all (*broadcast*) waiting threads

```
int pthread_cond_signal(pthread_cond_t *cond);
int pthread_cond_broadcast(pthread_cond_t *cond);
```

## Implementing Join

Pthread operations are the wrapper versions

```
int done = 0;
pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t c = PTHREAD_COND_INITIALIZER;

void thr_exit() {
    pthread_mutex_lock(&m);
    done = 1;
    pthread_cond_signal(&c);
    pthread_mutex_unlock(&m);
}

void *child(void *arg) {
    printf("child\n");
    thr_exit();
    return NULL;
}

void thr_join() {
    pthread_mutex_lock(&m);
    while (done == 0)
        pthread_cond_wait(&c, &m);
    pthread_mutex_unlock(&m);
}

int main(int argc, char *argv[]) {
    printf("parent: begin\n");
    pthread_t p;
    pthread_create(&p, NULL, child, NULL);
    thr_join();
    printf("parent: end\n");
    return 0;
}
```

have a state variable, not just wait and signal

it is good practice to hold the lock when calling signal, even though signal doesn't explicitly require it like wait does

lock must be held when wait is called – wait atomically releases the lock and puts the calling thread to sleep, then re-acquires the lock before returning when thread is awakened

loop instead of if

## Incorrect

Pthread operations are the wrapper versions

```

pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t c = PTHREAD_COND_INITIALIZER;

void thr_exit() {
    Pthread_mutex_lock(&m);
    Pthread_cond_signal(&c);
    Pthread_mutex_unlock(&m);
}

void *child(void *arg) {
    printf("child\n");
    thr_exit();
    return NULL;
}

void thr_join() {
    Pthread_mutex_lock(&m);
    Pthread_cond_wait(&c, &m);
    Pthread_mutex_unlock(&m);
}

int main(int argc, char *argv[]) {
    printf("parent: begin\n");
    pthread_t p;
    Pthread_create(&p, NULL, child, NULL);
    thr_join();
    printf("parent: end\n");
    return 0;
}

```

no state variable

it is possible for the child to run immediately, before the parent calls thr\_join and wait

if the child calls signal before the parent calls wait, the signal is lost and the parent is left waiting forever

if the child calls signal before the parent calls wait, the parent knows because done is 1 and thus doesn't wait

```

void thr_exit() {
    Pthread_mutex_lock(&m);
    done = 1;
    Pthread_cond_signal(&c);
    Pthread_mutex_unlock(&m);
}

void thr_join() {
    Pthread_mutex_lock(&m);
    while (done == 0)
        Pthread_cond_wait(&c, &m);
    Pthread_mutex_unlock(&m);
}

```

<https://pages.cs.wisc.edu/~remzi/OSTEP/threads-cv.pdf> 66

## Incorrect

Pthread operations are the wrapper versions

```

int done = 0;
pthread_cond_t c = PTHREAD_COND_INITIALIZER;

void thr_exit() {
    done = 1;
    Pthread_cond_signal(&c);
}

void *child(void *arg) {
    printf("child\n");
    thr_exit();
    return NULL;
}

void thr_join() {
    if (done == 0)
        Pthread_cond_wait(&c);
}

int main(int argc, char *argv[]) {
    printf("parent: begin\n");
    pthread_t p;
    Pthread_create(&p, NULL, child, NULL);
    thr_join();
    printf("parent: end\n");
    return 0;
}

```

no lock

non-atomic test and act – if the child runs between the parent checking done and waiting, the child can call signal before the parent calls wait

<https://pages.cs.wisc.edu/~remzi/OSTEP/threads-cv.pdf> 67

## Producer/Consumer (Bounded Buffer)

- one or more *producer* threads add to a buffer
- one or more *consumer* threads remove from the buffer
- e.g. multithreaded web server
  - requests come in from multiple clients
  - multiple threads on the server satisfy requests
- e.g. Unix pipes
  - output comes from one process and is consumed as input by another
- the buffer is a shared resource, so race conditions must be avoided

CPS31: Operating Systems • Spring 2026 68

## The Idea of Producer/Consumer

```

void *producer(void *arg) {
    int i;
    int loops = (int) arg;
    for (i = 0; i < loops; i++) {
        put(i);
    }
}

void *consumer(void *arg) {
    while (1) {
        int tmp = get();
        printf("%d\n", tmp);
    }
}

int buffer;
int count = 0; // initially, empty

void put(int value) {
    assert(count == 0);
    count = 1;
    buffer = value;
}

int get() {
    assert(count == 1);
    count = 0;
    return buffer;
}

```

producer adds to the buffer

consumer removes from the buffer

buffer holds (only) a single value (for now)

can only add to the buffer when it is empty (assert verifies that assumption – will crash if violated)

can only remove from the buffer when it has something in it (assert verifies that assumption – will crash if violated)

CPS31: Operating Systems • Spring 2026 69

## First Try

```
void *producer(void *arg) {
    int i;
    int loops = (int) arg;
    for (i = 0; i < loops; i++) {
        put(i);
    }
}

void *consumer(void *arg) {
    while (1) {
        int tmp = get();
        printf("%d\n", tmp);
    }
}
```

```
int buffer;
int count = 0; // initially, empty

void put(int value) {
    assert(count == 0);
    count = 1;
    buffer = value;
}

int get() {
    assert(count == 1);
    count = 0;
    return buffer;
}
```

- does this work?
  - no, unless the put and get calls are interleaved exactly right – two puts in a row or two gets in a row will result in a violated assertion
  - also race conditions within put and get themselves

## Second Try

add a condition variable to wait for expected buffer status

```
int loops; // must initialize somewhere...
cond_t cond;
mutex_t mutex;

void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        pthread_mutex_lock(&mutex); // p1
        if (count == 1) // p2
            pthread_cond_wait(&cond, &mutex); // p3
        put(i); // p4
        pthread_cond_signal(&cond); // p5
        pthread_mutex_unlock(&mutex); // p6
    }
}

void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        pthread_mutex_lock(&mutex); // c1
        if (count == 0) // c2
            pthread_cond_wait(&cond, &mutex); // c3
        int tmp = get(); // c4
        pthread_cond_signal(&cond); // c5
        pthread_mutex_unlock(&mutex); // c6
        printf("%d\n", tmp);
    }
}
```

```
int buffer;
int count = 0; // initially, empty

void put(int value) {
    assert(count == 0);
    count = 1;
    buffer = value;
}

int get() {
    assert(count == 1);
    count = 0;
    return buffer;
}
```

- does this work?
  - with a single producer and a single consumer
    - lock means can't be signaled between checking count and waiting
    - once waiting, only the other thread can wake
    - only one thread can be waiting at a time
  - with more than one producer and/or consumer
    - assumption is that when wait returns, the buffer is ready for us

## Second Try

two consumers, one producer

T <sub>1</sub>	State	T <sub>2</sub>	State	T <sub>p</sub>	State	Count	Comment
c1	Run	Ready	Ready	Ready	0		
c2	Run	Ready	Ready	Ready	0		
c3	Sleep	Ready	Ready	Ready	0		Nothing to get
	Sleep	Ready	p1	Run	0		
	Sleep	Ready	p2	Run	0		
	Sleep	Ready	p4	Run	1		Buffer now full
	Ready	Ready	p5	Run	1		T <sub>c1</sub> awoken
	Ready	Ready	p6	Run	1		
	Ready	Ready	p1	Run	1		
	Ready	Ready	p2	Run	1		
	Ready	Ready	p3	Sleep	1		Buffer full; sleep
	Ready	c1	Run	Sleep	1		T <sub>c2</sub> sneaks in ...
	Ready	c2	Run	Sleep	1		
	Ready	c4	Run	Sleep	0		... and grabs data
	Ready	c5	Run	Ready	0		T <sub>p</sub> awoken
	Ready	c6	Run	Ready	0		
c4	Run	Ready	Ready	Ready	0		Oh oh! No data

```
int loops; // must initialize somewhere...
cond_t cond;
mutex_t mutex;

void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        pthread_mutex_lock(&mutex); // p1
        if (count == 1) // p2
            pthread_cond_wait(&cond, &mutex); // p3
        put(i); // p4
        pthread_cond_signal(&cond); // p5
        pthread_mutex_unlock(&mutex); // p6
    }
}

void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        pthread_mutex_lock(&mutex); // c1
        if (count == 0) // c2
            pthread_cond_wait(&cond, &mutex); // c3
        int tmp = get(); // c4
        pthread_cond_signal(&cond); // c5
        pthread_mutex_unlock(&mutex); // c6
        printf("%d\n", tmp);
    }
}
```

## Mesa vs Hoare Semantics

Mesa semantics – (how pretty much every system works)

- signal only wakes a thread
  - being signaled reflects a certain state of the world at that moment but there are no guarantees that state still holds when the thread finally runs

Hoare semantics – (nice but harder to implement)

- signal wakes a thread *and* guarantees that it runs immediately

### Third Try while instead of if

```
int loops; // must initialize somewhere...
cond_t cond;
mutex_t mutex;

void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex); // p1
        while (count == 1) // p2
            Pthread_cond_wait(&cond, &mutex); // p3
        put(i); // p4
        Pthread_cond_signal(&cond); // p5
        Pthread_mutex_unlock(&mutex); // p6
    }
}

void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex); // c1
        while (count == 0) // c2
            Pthread_cond_wait(&cond, &mutex); // c3
        int tmp = get(i); // c4
        Pthread_cond_signal(&cond); // c5
        Pthread_mutex_unlock(&mutex); // c6
        printf("%d\n", tmp);
    }
}
```

```
int buffer;
int count = 0; // initially, empty

void put(int value) {
    assert(count == 0);
    count = 1;
    buffer = value;
}

int get() {
    assert(count == 1);
    count = 0;
    return buffer;
}
```

#### • does this work?



- the problem of a buffer status change between waking and access is fixed
- but a single condition variable means that both producers and consumers wait, and thus a thread can be woken by another thread of the same kind rather than the opposite...and if it can't do its job, it goes back to sleep and no one else gets signaled



### Third Try

```
int loops; // must initialize somewhere...
cond_t cond;
mutex_t mutex;

void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex); // p1
        while (count == 1) // p2
            Pthread_cond_wait(&cond, &mutex); // p3
        put(i); // p4
        Pthread_cond_signal(&cond); // p5
        Pthread_mutex_unlock(&mutex); // p6
    }
}
```

#### two consumers, one producer

T <sub>c1</sub>	State	T <sub>c2</sub>	State	T <sub>p</sub>	State	Count	Comment
c1	Run		Ready		Ready	0	
c2	Run		Ready		Ready	0	
c3	Sleep		Ready		Ready	0	Nothing to get
	Sleep	c1	Run		Ready	0	
	Sleep	c2	Run		Ready	0	
	Sleep	c3	Sleep		Ready	0	Nothing to get
	Sleep		Sleep	p1	Run	0	
	Sleep		Sleep	p2	Run	0	
	Sleep		Sleep	p4	Run	1	Buffer now full
	Ready		Sleep	p5	Run	1	T <sub>c1</sub> awoken
	Ready		Sleep	p6	Run	1	
	Ready		Sleep	p1	Run	1	
	Ready		Sleep	p2	Run	1	
	Ready		Sleep	p3	Sleep	1	Must sleep (full)
c2	Run		Sleep		Sleep	1	Recheck condition
c4	Run		Sleep		Sleep	0	T <sub>c1</sub> grabs data
c5	Run		Ready		Sleep	0	Oops! Woke T <sub>c2</sub>
c6	Run		Ready		Sleep	0	
c1	Run		Ready		Sleep	0	
c2	Run		Ready		Sleep	0	
c3	Sleep		Ready		Sleep	0	Nothing to get
	Sleep	c2	Run		Sleep	0	
	Sleep	c3	Sleep		Sleep	0	Everyone asleep...

```
void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex); // c1
        while (count == 0) // c2
            Pthread_cond_wait(&cond, &mutex); // c3
        int tmp = get(i); // c4
        Pthread_cond_signal(&cond); // c5
        Pthread_mutex_unlock(&mutex); // c6
        printf("%d\n", tmp);
    }
}
```

### Fourth Try separate condition variables for empty buffer and full buffer

```
int loops; // must initialize somewhere...
cond_t empty, fill;
mutex_t mutex;

void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex);
        while (count == 1)
            Pthread_cond_wait(&empty, &mutex);
        put(i);
        Pthread_cond_signal(&fill);
        Pthread_mutex_unlock(&mutex);
    }
}

void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex);
        while (count == 0)
            Pthread_cond_wait(&fill, &mutex);
        int tmp = get(i);
        Pthread_cond_signal(&empty);
        Pthread_mutex_unlock(&mutex);
        printf("%d\n", tmp);
    }
}
```

```
int buffer;
int count = 0; // initially, empty

void put(int value) {
    assert(count == 0);
    count = 1;
    buffer = value;
}

int get() {
    assert(count == 1);
    count = 0;
    return buffer;
}
```

#### • does this work?

😊 yes!

### Full Producer/Consumer Solution

```
int buffer[MAX];
int fill_ptr = 0;
int use_ptr = 0;
int count = 0;

void put(int value) {
    buffer[fill_ptr] = value;
    fill_ptr = (fill_ptr + 1) % MAX;
    count++;
}

int get() {
    int tmp = buffer[use_ptr];
    use_ptr = (use_ptr + 1) % MAX;
    count--;
    return tmp;
}
```

```
int loops; // must initialize somewhere...
cond_t empty, fill;
mutex_t mutex;

void *producer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex);
        while (count == MAX)
            Pthread_cond_wait(&empty, &mutex);
        put(i);
        Pthread_cond_signal(&fill);
        Pthread_mutex_unlock(&mutex);
    }
}

void *consumer(void *arg) {
    int i;
    for (i = 0; i < loops; i++) {
        Pthread_mutex_lock(&mutex);
        while (count == 0)
            Pthread_cond_wait(&fill, &mutex);
        int tmp = get(i);
        Pthread_cond_signal(&empty);
        Pthread_mutex_unlock(&mutex);
        printf("%d\n", tmp);
    }
}
```

add more slots to the buffer

use a circular array implementation of a queue – use\_ptr at the head, fill\_ptr at the tail – to avoid shifting and achieve O(1) operations

not thread safe – why not, and why is this not a problem?